

PCTWORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

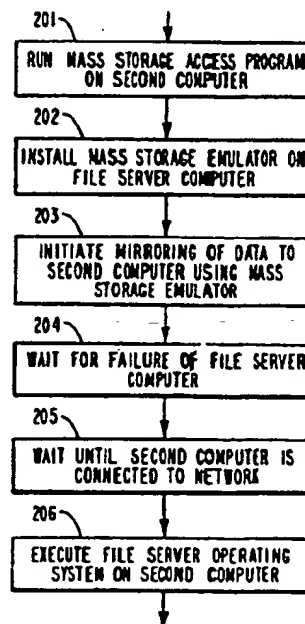
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁵ : G06F 11/20, 11/16		A1	(11) International Publication Number: WO 95/03580 (43) International Publication Date: 2 February 1995 (02.02.95)
(21) International Application Number: PCT/US94/07046 (22) International Filing Date: 21 June 1994 (21.06.94) (30) Priority Data: 08/094,755 20 July 1993 (20.07.93) US (71) Applicant: VINCA CORPORATION [US/US]; 4000 Central Park East, 1815 South State Street, Orem, UT 84058 (US). (72) Inventors: OHRAN, Richard; 71 West 4750 North, Provo, UT 84604 (US). DICKSON, Terry; 9233 Country Glen Circle, Sandy, UT 84093 (US). (74) Agents: CHRISTIANSEN, Jon, C. et al.; Vancott, Bagley, Cornwall & McCarthy, Suite 1600, 50 South Main Street, Salt Lake City, UT 84144 (US).			(81) Designated States: AT, AU, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, ES, FI, GB, HU, JP, KP, KR, KZ, LK, LU, LV, MG, MN, MW, NL, NO, NZ, PL, PT, RO, RU, SD, SE, SK, UA, UZ, VN, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG). Published With international search report.

(54) Title: METHOD FOR RAPID RECOVERY FROM A NETWORK FILE SERVER FAILURE

(57) Abstract

A method for providing rapid recovery from a network file server failure through the use of a backup computer system. Unlike other redundant file server configurations, this method does not require the backup computer system to be running the file server operating system. Instead, it runs a special mass storage access program (201) that communicates with a mass storage emulator program on the network file server (202), making the disks (or other mass storage devices) on the backup computer system appear like they were disks on the file server computer. By mirroring data by writing to both the mass storage of the file server and through the mass storage emulator and mass storage access program to the disks on the backup computer, a copy of the data on the file server computer is made (203). In the event of failure of the file server computer (204), the backup computer can be restarted as a file server (205), using the copy on its disks. The same method can be utilized to restore normal system operation when the failure has been corrected.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

1 METHOD FOR RAPID RECOVERY FROM A
2 NETWORK FILE SERVER FAILURE

3 Microfiche Appendix

4 This specification includes a MICROFICHE
5 APPENDIX which is intended to be a part of the
6 disclosure herein. The MICROFICHE APPENDIX
7 contains 1 page of microfiche and a total of 69
8 frames. The MICROFICHE APPENDIX includes computer
9 source code used in the implementation of one
10 preferred embodiment of the invention.

11 Background of the Invention

12 Field of the Invention. This invention relates to
13 network file server computer systems, and in
14 particular to the methods used to recover from a
15 computer failure in a system with a plurality of
16 computer systems, each with its own mass storage
17 devices.

18 Description of Related Art. It is often desirable
19 to provide continuous operation of computer

1 systems, particularly file servers which support a
2 number of user workstations or personal computers.
3 To achieve this continuous operation, it is
4 necessary for the computer system to be tolerant of
5 software and hardware problems or faults. This is
6 generally done by having redundant computers and
7 mass storage devices, such that a backup computer
8 or disk drive is immediately available to take over
9 in the event of a fault.

10 A number of techniques for implementing a fault-
11 tolerant computer system are described in Major et
12 al., United States Patent 5,157,663 (which is
13 hereby incorporated by reference in its entirety)
14 and its cited references. In particular, the
15 invention of Major provides a redundant network
16 file server capable of recovering from the failure
17 of either the computer or the mass storage device
18 of one of the file servers. The file server
19 operating system is run on each computer system in
20 the network file server, with each computer system
21 cooperating to produce the redundant network file
22 server. This technique has been used by Novell to

1 implement its SFT-III fault-tolerant file server
2 product.

3 There are a number of reasons why the use of a
4 redundant network file server such as that
5 described in Major may be undesirable. As can be
6 seen from the description in Major, the software
7 needed to provide such a redundant network file
8 server is considerably more complex than the
9 software of the present invention. This can result
10 in a lower reliability due the increased presence
11 of programming errors ("bugs") in the complex
12 software. Also, the processing time required to
13 handle a client request may be increased by the
14 complexity of the redundant network file server
15 software, when compared to a single-processor
16 network file server. Finally, license restrictions
17 or other limitations may make it infeasible or
18 uneconomical to run a redundant network file server
19 instead of a normal network file server.

20 Summary of the Invention

1 It is an object of this invention to provide
2 rapid recovery from a network file server failure
3 without the complex software of a redundant network
4 file server. This is achieved by having a second,
5 backup computer system with its own mass storage
6 device (generally a magnetic disk). This backup
7 computer is connected by an appropriate means for
8 communications to the file server computer,
9 allowing the transmission of information (such as
10 commands and data) between the two computers. A
11 mass storage emulator, running like a device driver
12 on the file server computer, sends information to a
13 mass storage access program on the backup computer.
14 The mass storage access program performs the
15 requested operation (read, write, etc.) on the mass
16 storage system connected to the backup computer,
17 and returns the result to the mass storage emulator
18 on the file server computer.

19 This makes the mass storage device on the backup
20 computer look like another mass storage device on
21 the file server computer. The data mirroring
22 option of the file server operating system can be

1 activated (or, if the operating system does not
2 support data mirroring, a special device driver
3 that provides data mirroring can be used), so that
4 a copy of all data written to the mass storage
5 device directly connected to the file server will
6 also be written to the mass storage device on the
7 backup computer, through the mass storage emulator
8 and mass storage access programs.

9 When a failure is detected in the file server
10 computer system, the backup computer can be
11 restarted as a file server computer. The mass
12 storage device of the backup computer will contain
13 a copy of the information on the mass storage
14 device of the failed file server, so the new file
15 server can start with approximately the same data
16 as when the previous file server failed.

17 It is a further object of this invention to
18 improve the reliability of a redundant network file
19 server computer system by reducing the complexity
20 of the software when compared to the software of a
21 redundant network file server. As will be clear
22 from the discussion of the preferred embodiment of

1 the invention, the programs for the mass storage
2 emulator on the file server computer and the mass
3 storage access on the backup computer are
4 considerably less complex than a full redundant
5 file server operating system.

6 It is a further object of this invention to
7 improve the performance of a network file server by
8 improving the performance of the mass storage used
9 by the network file server. Because of the
10 simplicity of the mass storage emulator and the
11 mass storage access program, as well as the fact
12 that the backup computer is only running the mass
13 storage access program and not a file server
14 operating system that must also handle network
15 requests and other activity, the performance of the
16 emulated mass storage device may exceed the
17 performance of a mass storage device directly
18 attached to the file server computer. This is
19 particularly true if the mass storage access
20 program is expanded to provide caching of
21 information and the full memory of the backup
22 computer (less that occupied by a simple operating

1 system and the mass storage access program) is used
2 as a cache.

3 These and other features of the invention will
4 be more readily understood upon consideration of
5 the attached drawings and of the following detailed
6 description of those drawings and the presently
7 preferred embodiments of the invention.

8 Brief Description of the Drawings

9 Figure 1 illustrates a computer configuration on
10 which the method of the invention runs.

11 Figure 2 is a flow diagram showing the steps of
12 the method of the invention.

13 Detailed Description of the Invention

14 Referring to Figure 1, which illustrates a
15 representative computer configuration on which the
16 method of the invention runs, it can be seen that
17 there are two computer systems 110 and 120. The
18 first computer system 110 is running a file server
19 operating system (such as Novell NetWare®).
20 Computer system 110 includes computer 112 connected

1 to network 101 through interface 111 (and its
2 associated software), and mass storage device 114
3 connected through controller 113 (and its
4 associated software). These represent the standard
5 components of a network file server. In the case
6 of NetWare, computer 112 is generally a PC-
7 compatible computer based on an Intel 386 or 486
8 processor, network 101 can be an ethernet (so that
9 interface 111 is an ethernet interface), and mass
10 storage device 114 is an SCSI or IDE magnetic disk
11 connected through an appropriate controller 113.
12 Computer 122 would also be a PC-compatible
13 computer, so that it could also run the same
14 NetWare file server operating system as computer
15 112. Network 101 could also be implemented as a
16 token ring, Arcnet, or any other network
17 technology, such network technology being known to
18 those skilled in the art.

19 The mass storage devices of the invention should
20 not be viewed as limited to magnetic disk drives,
21 but can also be implemented using optical discs,
22 magnetic tape drives, or any other medium capable

1 of handling the read and write requests of the
2 particular computer system.

3 Added to the standard network file server to
4 support the method of this invention are a backup
5 computer system 120 and a means 102 for
6 communicating between computer system 110 and
7 computer system 120.

8 Computer system 120 has components similar to
9 computer system 110. Computer system 120 can be
10 connected to network 101 through interface 121,
11 although it is not necessary for computer system
12 120 to actually be connected to network 101 during
13 normal operation. Computer 122 is connected to
14 interface 121 and to mass storage device 124
15 through controller 123.

16 While it is not necessary for computer system
17 120 to have identical components to computer system
18 110, in many cases that will be the case. In other
19 cases, computer system 120 may be an older, slower
20 system previously used as a file server but
21 replaced with computer system 110. All that is
22 required of computer system 120 is that it be

1 capable of running the file server operating system
2 in case of the failure of computer system 110, and
3 that its mass storage device 124 be of sufficient
4 capacity to hold the data mirrored from mass
5 storage device 114.

6 Communications means 102 provides a link between
7 computer systems 110 and 120. Computer 112 is
8 connected to communications means 102 through
9 attachment 115, and computer 122 is connected to
10 communications means 102 through attachment 125.
11 Communications means 102 can be implemented using a
12 variety of techniques, well-known to those skilled
13 in the art. In the preferred embodiment, a high-
14 speed serial point-to-point link is used. An
15 alternative would be to use the serial
16 communications ports of computers 112 and 122,
17 programmed to run at a high data rate, or the
18 parallel interfaces of computers 112 and 122.
19 Another alternative is for communications means 102
20 to be a virtual circuit or channel carried on
21 network 101. In this latter case, communications
22 means 102 would really be network 101, attachment

1 115 would really be interface 111, and attachment
2 125 would really be interface 121.

3 It is important that communication means 102
4 provide data transfer at rates comparable to the
5 data rate of mass storage device 124 so that it
6 does not limit the performance of the system. The
7 method of this invention is not dependant on the
8 particular implementation of communications means
9 102.

10 Figure 2 is a flow diagram showing the steps of
11 the method of the invention. In step 201, a
12 special program -- the mass storage access program
13 -- is run on computer system 120. The mass storage
14 access program receives commands from computer
15 system 110 over communications means 102. Based on
16 those commands, the mass storage access program
17 accesses mass storage device 124 to perform the
18 operation specified in the command received from
19 computer system 110. The results of the accessing
20 of mass storage device 124 is returned to computer
21 system 110 over communications means 102.

1 The mass storage access program can be enhanced
2 to provide a cache of data on mass storage device
3 124. The implementation of such a cache function
4 is well-known in the art, consisting of keeping a
5 copy of the most recently accessed information of
6 mass storage device 124 in the memory of computer
7 122. When a read command is received, it is not
8 necessary to access mass storage device 124 if a
9 copy of the data is in the cache. Since, in the
10 preferred embodiment, computer 122 has a large
11 memory (it must be large enough to run the file
12 server operating system) and the mass storage
13 access program is quite small, there is a large
14 amount of memory available for the cache. This
15 means that many entries will be in the cache, and
16 the chance of finding a block being read in the
17 cache is higher than would be normal for a similar
18 cache in a file server operating system.

19 In step 202, coincidentally with the running of
20 the mass storage access program on computer system
21 120, another program -- the mass storage emulator -
22 - is installed on computer system 110. The mass

1 storage emulator takes mass storage requests from
2 the file server operating system running on
3 computer system 110 and sends them as commands over
4 communications means 102 to computer system 120,
5 where they are processed by the mass storage access
6 program, as discussed above.

7 When results from a command are received from
8 the mass storage access program over communications
9 means 102 by the mass storage emulator, they are
10 returned to the file server operating system, much
11 as the result of a normal mass storage request
12 would be returned. In this way, the mass storage
13 access program and the mass storage emulator
14 cooperate to make it appear to the file server
15 operating system that mass storage device 124 is
16 directly connected to computer 112 on computer
17 system 110.

18 In the preferred embodiment of this invention,
19 the mass storage access program is a conventional
20 program running under the disk operating system of
21 personal computer 122. The disk storage emulator
22 is a NetWare Loadable Module (NLM), much like the

1 device driver for a disk drive. Copies of the
2 source code for the mass storage access program and
3 the mass storage emulator are given in the
4 microfiche appendix.

5 In step 203, mirroring of data is initiated.
6 When data is being mirrored on two or more mass
7 storage devices, whenever data is to be written it
8 is written to all mass storage devices taking part
9 in the mirroring, at the same location on each mass
10 storage device. (The location may be relative to
11 the start of the mass storage device, or to the
12 start of a partition or contiguous portion of the
13 mass storage device, as appropriate to the way the
14 mass storage device has been formatted and is being
15 used.) Data can be read from any mass storage
16 device taking part in the mirroring, since each
17 mass storage device contains identical data.

18 Mirroring may be an integral function of the
19 file server operating system, so that no special
20 program is necessary for implementing disk
21 mirroring as part of the method of this invention.
22 Step 203 only requires the activation or starting

1 of mirroring on the part of the file server
2 operating system. This is the case in the
3 preferred embodiment of the invention, operating
4 with NetWare and using the mirroring facilities of
5 that file server operating system.

6 If the file server operating system does not
7 provide mirroring, a separate mirroring module will
8 have to be implemented. Such a mirroring module,
9 whose implementation should be obvious to one
10 skilled in the art, will take each write request
11 and pass it to the driver for each mass storage
12 device taking part in the mirroring. For mass
13 storage device 124 on computer system 120, the
14 driver will be the mass storage emulator, discussed
15 above. When successful completion of the write
16 request has been received from all mass storage
17 devices taking part in the mirroring, the mirroring
18 module will indicate successful completion to the
19 file server operating system.

20 For read requests, the mirroring module can
21 direct the read request to any of the mass storage
22 devices, since all contain identical data.

1 Generally, the read request will be directed to the
2 mass storage device which is first available to
3 handle the request.

4 As part of the initiating of mirroring, it is
5 necessary to assure that each mass storage device
6 taking part in mirroring has the same contents at
7 the start of mirroring. This can be done by
8 designating one of the mass storage devices as the
9 master, and making a copy of the master mass
10 storage device's data to all other mass storage
11 devices taking part in the mirroring. An
12 alternative approach is to have a timestamp
13 indicating when the last change was made to the
14 data on a mass storage device. If the timestamp on
15 a mass storage device is the same as the timestamp
16 on the master mass storage device, it will not be
17 necessary to make a new copy of the data.

18 At step 204, the method of this invention waits
19 until a failure of file server computer system 110
20 is detected. Such a failure could come from the
21 failure of either hardware (such as computer 112 or
22 mass storage device 114) or software (such as the

1 file server operating system). Although means for
2 automatically detecting such a failure may be used,
3 in the preferred embodiment such failure is
4 detected by a system operator or workstation user
5 by noticing that file server requests are no longer
6 being handled by computer system 110. It is not
7 difficult for a user to determine there is a
8 problem with file server computer system 110; in
9 most cases, a user workstation will stop working
10 and "hang" while it waits for a file server request
11 that will never be honored.

12 In step 205, when a failure of computer system
13 110 has been detected, if computer system 120 is
14 not currently connected to network 101 through
15 interface 121, a connection to the network 101 is
16 made. This can be done either by activating
17 interface 121 or physically connecting interface
18 121 to network 101, as appropriate.

19 In step 206, when computer system 120 has been
20 connected to network 101, a file server operating
21 system is loaded into computer 122 and executed,
22 making computer system 120 a file server computer

1 system. New file server computer system 120 now
2 responds to requests received from network 101 as
3 failed file server computer system 110 did before
4 its failure. The file server operating system
5 executing on computer 122 accesses mass storage
6 device 124 to respond to the requests.

7 Note that because mass storage device 124
8 received data through the mass storage emulator and
9 mass storage access program while file server
10 computer system 110 was operating, mass storage
11 device 124 contains a copy of the data stored on
12 mass storage device 114 prior to the failure of
13 computer system 120. (Because of timing, the last
14 few write operations may not have occurred on all
15 mass storage devices taking part in mirroring, but
16 the file server operating system is capable of
17 handling these small differences.) Because a copy
18 of the mass storage data of failed file server
19 computer system 110 is immediately available to new
20 file server computer system 120, the time necessary
21 to recover from a file server failure is minimal.

1 When the fault that caused the failure of
2 computer system 120 has been corrected, fault-
3 tolerant operation can be restored. Depending on
4 the relative capabilities of computer systems 110
5 and 120, one of two techniques can be employed.
6 Both involve the same method steps as were
7 discussed above.

8 If the two computer systems have components of
9 similar speed and capacity, there is no reason not
10 to continue using computer system 120 as the file
11 server computer. In this case, computer system 110
12 can now be treated as the backup computer system.
13 The mass storage access program is run on computer
14 system 110, the mass storage emulator is installed
15 on computer system 120, and mirroring is initiated
16 on the file server operating system running on
17 computer system 120. As part of the initiating of
18 mirroring, any data written to mass storage device
19 124 during the time computer system 110 was not
20 available is now copied to mass storage device 114
21 though the mass storage emulator, communications
22 mean 102, and the mass storage access program.

1 Alternatively, if computer system 120 is less
2 capable than computer system 110, it will be
3 desirable to make computer system 110 the file
4 server computer system when the failure has been
5 corrected. To accomplish this, two approaches are
6 possible. In the first approach, computer system
7 110 is brought up as the backup computer system,
8 running the mass storage access program, as
9 discussed above. When mass storage device 114
10 contains a copy of the data on mass storage device
11 124, computer system 110 can be restarted as the
12 file server (running the file server operating
13 system) and computer system 120 can be restarted as
14 the backup computer in accordance with the method
15 discussed above.

16 Alternatively, when the failure of computer
17 system 110 has been corrected, computer system 120
18 can be restarted as backup computer system, running
19 the mass storage access program, and computer
20 system 110 can be restarted as the file server
21 computer, running the file server operating system
22 and the mass storage emulator. When mirroring is

1 initiated, it will be determined by the timestamps
2 stored on each of mass storage devices 114 and 124
3 that the data on mass storage device 114 is out of
4 date. The file server operating system will read
5 the data on mass storage device 124 (though the
6 mass storage emulator, communications means 102,
7 and the mass storage access program). It will also
8 copy the data from mass storage device 124 to mass
9 storage device 114 until they contain identical
10 data.

11 It is to be understood that the above described
12 embodiments are merely illustrative of numerous and
13 varied other embodiments which may constitute
14 applications of the principles of the invention.
15 Such other embodiments may be readily devised by
16 those skilled in the art without departing from the
17 spirit or scope of this invention and it is our
18 intent they be deemed within the scope of our
19 invention.

1 Claims

2 We claim:

3 1. A method for rapid recovery from a network file
4 server failure, operating on a computer
5 configuration that includes:

6 a first computer system, comprising:

7 (a) a first computer executing a file
8 server operating system adapted to
9 respond to requests received from a
10 network,

11 (b) a first interface connecting said
12 first computer to said network, and

13 (c) a first mass storage device connected
14 to said first computer;

15 a second computer system, comprising:

16 (a) a second computer capable of executing
17 said file server operating system,

18 (b) a second interface capable of
19 connecting said second computer to
20 said network, and

21 (c) a second mass storage device connected
22 to said second computer;

1 and means for communicating between said first
2 computer system and said second computer system,
3 the recovery method of the invention comprising:
4 (A) running a mass storage access program on
5 said second computer, said mass storage
6 access program receiving commands from said
7 first computer over said communicating
8 means, accessing said second mass storage
9 device as specified by said commands, and
10 returning said commands' results to said
11 first computer over said communicating
12 means;
13 (B) installing a mass storage emulator on said
14 first computer for use by said file server
15 operating system, said mass storage
16 emulator taking mass storage requests from
17 said file server operating system, sending
18 commands to said second computer system
19 over said communicating means, and
20 returning results of said commands received
21 over said communicating means from said

1 second computer to said file server
2 operating system;
3 (C) initiating mirroring of data by writing
4 data both to said first mass storage device
5 and through said mass storage emulator and
6 said mass storage access program to said
7 second mass storage device;
8 (D) when a failure of said first computer
9 system is detected, stopping said mass
10 storage access program on said second
11 computer; and then
12 (E) if said second computer is not currently
13 connected to said network, connecting said
14 second computer to said network through
15 said second interface; and then
16 (F) executing said file server operating system
17 on said second computer, said file server
18 operating system responding to requests
19 received from said network and accessing
20 data stored on said second mass storage
21 device.

- 1 2. A method as in claim 1, wherein said first
2 computer system and said second computer system are
3 each PC-compatible computers.
- 4 3. A method as in claim 2 wherein said file server
5 operating system is Novell NetWare.
- 6 4. A method as in claim 1, wherein said network is
7 an ethernet.
- 8 5. A method as in claim 1, wherein said network is
9 a token ring.
- 10 6. A method as in claim 1, wherein said first mass
11 storage device is a magnetic disk.
- 12 7. A method as in claim 1, wherein said second
13 mass storage device is a magnetic disk.
- 14 8. A method as in claim 1, wherein said means for
15 communicating between said first computer system
16 and said second computer system is a high-speed
17 serial point-to-point link.
- 18 9. A method as in claim 1, wherein said means for
19 communicating between said first computer system
20 and said second computer system uses the serial
21 communications ports of said first computer and
22 said second computer.

- 1 10. A method as in claim 1, wherein said means for
2 communicating between said first computer system
3 and said second computer system uses the parallel
4 interfaces of said first computer and said second
5 computer.
- 6 11. A method as in claim 1, wherein said means for
7 communicating between said first computer system
8 and said second computer system is said network.
- 9 12. A method as in claim 1, wherein said mass
10 storage access program provides a cache of data on
11 said second mass storage device.
- 12 13. A method as in claim 1, wherein said mass
13 storage access program and said mass storage
14 emulator makes said second mass storage device of
15 said second computer system look like another mass
16 storage device on said first computer system.
- 17 14. A method as in claim 1, wherein said mirroring
18 of data is performed by said file server operating
19 system.
- 20 15. A method as in claim 1, wherein said mirroring
21 of data is performed by a mirroring module.

1 16. A method as in claim 1, wherein said method is
2 also used for restoration of fault-tolerant
3 operation.

4 17. A method as in claim 16, wherein said second
5 computer system becomes a file server computer
6 system and said first computer system becomes a
7 backup computer system.

8 18. A method as in claim 16, wherein said first
9 computer system again becomes a file server
10 computer and said second computer system becomes a
11 backup computer system.

12 19. A method for rapid recovery from a network
13 file server failure, the method operating on a
14 computer system that includes:

15 a first computer system, comprising:

16 (a) a first computer capable of executing
17 a file server operating system adapted
18 to respond to requests received from a
19 network,

20 (b) a first interface adapted to connect
21 said first computer to said network,
22 and

1 (c) a first mass storage device capable of
2 being connected to said first
3 computer;
4 a second computer system, comprising:
5 (a) a second computer capable of executing
6 said file server operating system,
7 (b) a second interface capable of
8 connecting said second computer to
9 said network, and
10 (c) a second mass storage device capable
11 of being connected to said second
12 computer;
13 and means for communicating between said first
14 computer system and said second computer system,
15 the recovery method of the invention comprising:
16 (A) running a mass storage access program on
17 said second computer, said mass storage
18 access program being capable of receiving
19 commands from said first computer over said
20 communicating means, accessing said second
21 mass storage device as specified by said
22 commands, and returning said commands'

1 results to said first computer over said
2 communicating means;
3 (B) installing a mass storage emulator on said
4 first computer for use by said file server
5 operating system, said mass storage
6 emulator taking mass storage requests from
7 said file server operating system, sending
8 commands to said second computer system
9 over said communicating means, and
10 returning results of said commands received
11 over said communicating means from said
12 second computer to said file server
13 operating system;
14 (C) initiating mirroring of data by writing
15 data both to said first mass storage device
16 and through said mass storage emulator and
17 said mass storage access program to said
18 second mass storage device;
19 (D) when a failure of said first computer
20 system is detected, stopping said mass
21 storage access program on said second
22 computer;

- 1 (E) connecting said second computer to said
2 network through said second interface as
3 necessary;
- 4 (F) executing said file server operating system
5 on said second computer, said file server
6 operating system being capable of
7 responding to requests received from said
8 network and accessing data stored on said
9 second mass storage device.




FIG. 1

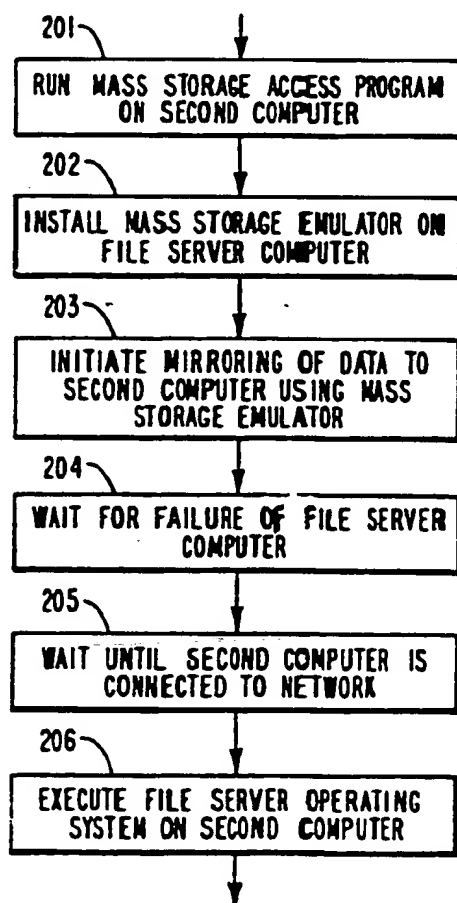


FIG. 2

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US94/07046

A. CLASSIFICATION OF SUBJECT MATTER

IPC(5) :G06F 11/20,16

US CL :395/575,600; 371/9.1,11.3

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 395/575,600; 371/9.1,11.3

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

PROQUEST (IEEE PERIODICALS)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y,P	US, A, 5,307,481 (Shimazaki et al) 26 April 1994 see Figures 1-3 and col. 2, lines 15-30; col. 5, line 22 - col. 8, line 19	1-19
Y	US, A, 5,157,663 (Major et al) 20 October 1992 see Figure 1; col. 2, line 42 - col. 3, line 32; col. 3, line 67 - col. 6, line 16; col. 9, lines 15 - 36	1-19
Y,P	US, A, 5,235,700 (Alaiwan et al) 10 August 1993 see Figure 1 and col. 4, line - col. 6, line 35	1-19
A	US, A, 5,079,740 (Patel et al) 07 January 1992, see Figure 1 and col. 1, line 46 - col. 2, line 44.	1-19



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:	*T	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be part of particular relevance	*X*	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	*Y*	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*A*	document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means		
P document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search

22 August 1994

Date of mailing of the international search report

26 SEP 1994

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. NOT APPLICABLE

Authorized officer

Paul Harrity

Telephone No. (703) 305-9677

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US94/07046

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US, A, 4,823,256 (Bishop et al) 18 April 1989, see Figures 1 and 3b, col. 2, line 39 - col. 3, line 44.	1-19
A, P	US, A, 5,295,258 (Jewett et al) 15 March 1994, see col. 2, line 38 - col. 4, line 61.	1-19